

TORIC ALGEBRA OF HYPERGRAPHS

SONJA PETROVIĆ AND DESPINA STASI

ABSTRACT. The edges of any hypergraph parametrize a monomial algebra called the edge subring of the hypergraph. We study presentation ideals of these edge subrings, and describe their generators in terms of balanced walks on hypergraphs. Our results generalize those for the defining ideals of edge subrings of graphs, which are well-known in the commutative algebra community, and popular in the algebraic statistics community. One of the motivations for studying toric ideals of hypergraphs comes from algebraic statistics, where generators of the toric ideal give a basis for random walks on fibers of the statistical model specified by the hypergraph. Further, understanding the structure of the generators gives insight into the model geometry.

1. INTRODUCTION

The rich history of the ideal theory of graphs, dating back to the seminal paper [SVV94], shows that fundamental properties of monomial algebras associated to graphs have interesting combinatorial interpretations. The connection inspired a body of work that joins combinatorial and algebraic methods with at least two algebraic objects in focus: monomial edge ideals of graphs, and toric ideals of edge subrings of graphs.

The study of quadratic squarefree monomial ideals corresponding to edges of a graph has generated a lot of interest from the algebraic community. In particular, the comprehensive monograph [Vil00] contains numerous references on the subject. More recent work includes, for example, [CN09] who describe Betti numbers and primary decompositions for edge ideals of restricted families of graphs. On the other hand, presentation ideals of edge subrings of graphs, which are toric rather than monomial, arise as kernels of monomial maps whose images correspond to the edges of a graph. These toric ideals have been studied in [Vil95], [OH99, OH00], [Vil00], [RTT12].

Several results for the monomial edge ideals have been extended to the case of hypergraphs (see, for example, [HT08], [DV11]), where a hypergraph is a generalization of a graph with edges containing more than two vertices. However, toric ideals of edge subrings remain completely unexplored beyond the graph case. Nevertheless, they are of interest in applications, and as blow-up algebras, since the edge subring of a uniform hypergraph is isomorphic to the special fiber ring of its edge ideal.

In this work, we extend the theory for these toric ideals from graphs to (uniform) hypergraphs. We study the combinatorial signatures of presentation ideals of monomial subalgebras parametrized by squarefree monomials of degree greater than 2. Their generators, Gröbner and Graver bases are given in terms of monomial walks on hypergraphs, generalizing the constructions that are well-known for graphs.

One of the motivations for studying toric ideals of hypergraphs comes from algebraic statistics, where generating sets of toric ideals play an important role in testing how well a model fits the given data. The edge subring of a hypergraph corresponds to any exponential family model whose joint probabilities are parametrized by monomials. For more details, see Section 3. In fact, edge subrings of graphs are already known in algebraic statistics, most notably in [Mor13], [OHT13], [SW12], [PRF10]. The first author has used these toric ideals to gain insight into the geometry of a particular random graph model, existence of maximum likelihood estimators, and a better way to generate Markov moves [PRF10, RPF]. In particular, this work is a starting point for developing the theory required for the more general algebraic statistical models.

A hypergraph H is d -uniform if every edge contains d vertices. We describe the natural one-to-one correspondence between sets of squarefree monomials of degree d and d -uniform hypergraphs with the following notation. For a finite d -uniform hypergraph H on $V = \{x_1, \dots, x_n\}$ and edge set E , each edge e_i of H encodes a squarefree monomial $x^{e_i} := \prod_{j \in e_i} x_j$ of degree d in the polynomial ring $k[x_1, \dots, x_n]$ over a field k . Thus,

the hypergraph H can be written as a set of monomials representing the edges: $H = \{x^{e_1}, \dots, x^{e_k}\}$. The *edge subring* of the hypergraph H , denoted by $k[H]$, is the monomial subring of the polynomial ring generated by the edges of H ; in symbols,

$$k[H] := k[x^{e_i} : e_i \in E(H)].$$

Letting t_{e_i} be a variable representing the edge e_i , define a ring homomorphism $\phi_H : k[t_{e_i}] \rightarrow k[H]$ with $\phi_H(t_{e_i}) = x^{e_i}$. The kernel of this monomial map, denoted by I_H , is the *toric ideal of the edge subring of the hypergraph H* . It encodes the algebraic relations among the edges of the hypergraph.

A first problem of interest is to describe the combinatorics of generators of the ideal for an arbitrary hypergraph. With this in mind, we define *monomial walks* on hypergraphs and the balancing condition in Section 2, which are the natural direct generalization of monomial walks on graphs. Theorem 2.8 characterizes the binomials in I_H in terms of such monomial walks, extending the classical theorems of Villarreal [Vil95, Vil00], and Ohsugi and Hibi [OH99]. In Section 3 we outline the relevance of our results: the correspondence between statistical models and hypergraphs is explored in Section 3.1, and we derive a connection to the well-studied set-theoretical problem of combinatorial discrepancy in Section 3.2. Given the apparent difficulty of the problem of describing *all* primitive monomial walks on an arbitrary hypergraph, we dedicate Section 4 to sparsely intersecting hypergraphs, which generalize the three types of primitive walks on graphs. In particular, Propositions 4.2, 4.5, 4.9 and Theorem 4.12 describe several supporting hypergraphs of primitive monomial walks. Their basic building blocks consist of matchings and sunflowers, well-known in the hypergraph literature. In Section 5 we show that many of the results for graphs cannot be generalized, and the expected degree bounds do not hold. Finally, Section 6 outlines some of the many open problems that we hope will be addressed by the combinatorics community.

Algebraic properties of toric ideals of hypergraphs such as normality, Cohen-Macaulayness, and geometry of the corresponding polytopes are all open problems of interest, but are beyond the scope of this paper. Degree bounds of generators are studied in [GP]. The problem of finding these generators algorithmically and understanding the complexity of such algorithms remains open and would make a significant contribution to applied algebraic statistics.

2. MONOMIAL WALKS ON HYPERGRAPHS

We begin with an example in the case when H is 2-uniform.

Example 2.1. Consider the simplest case when H is a graph, say, the complete graph $H = K_5$ on vertices x_1, \dots, x_5 . The edge subring $k[K_5] \subseteq k[x_1, \dots, x_5]$ is parametrized by

$$\begin{aligned} \phi_{K_5} : k[t_{ij}] &\rightarrow k[x_1x_2, x_1x_3, \dots, x_4x_5] \\ t_{ij} &\mapsto x_ix_j. \end{aligned}$$

The toric ideal I_{K_5} is generated by the following 10 quadrics:

$$\begin{aligned} t_{24}t_{35} - t_{23}t_{45}, \quad t_{14}t_{35} - t_{13}t_{45}, \quad t_{34}t_{25} - t_{23}t_{45}, \quad t_{14}t_{25} - t_{12}t_{45}, \quad t_{13}t_{25} - t_{12}t_{35}, \\ t_{34}t_{15} - t_{13}t_{45}, \quad t_{24}t_{15} - t_{12}t_{45}, \quad t_{23}t_{15} - t_{12}t_{35}, \quad t_{13}t_{24} - t_{12}t_{34}, \quad t_{23}t_{14} - t_{12}t_{34}. \end{aligned}$$

For example, the binomial $f := t_{13}t_{24} - t_{12}t_{34} \in I_H$ is represented in the leftmost graph of Figure 1. Since f is represented by an even cycle with alternating colors on the edges, we say that it *arises* from the cycle.

In general, when H is a graph, then the generators of the ideal I_H are known in terms of these red-blue colorings on the edges of the graph:

Theorem 2.2 ([Vil95], [OH99], see also [Vil00]). *The toric ideal of the edge subring of a graph G is generated by binomials arising from (primitive) even closed walks on G .*

Such binomials have been characterized by the following result which, in fact, provides a Graver basis for the toric ideal.

Theorem 2.3 ([Vil95], [OH99], see also [Vil00]). *Primitive even closed walks on G are one of the following: (i) even cycles, (ii) two odd cycles sharing a vertex, or (iii) two odd cycles such that there are two walks connecting a vertex in one cycle with a vertex in the other.*




FIGURE 1. The three types of walks from Theorem 2.3. Colors are explained in Remark 2.5.

An even closed walk is the exact structure in a graph required to give rise to a binomial: the sequence of edges in the walk induces a bipartition of the edges with multiplicity and all vertices in the walk have the same degree in each of the two parts. We propose the following generalization of an even closed walk on a hypergraph. Following [Vil00], we call it a monomial walk.

Definition 2.4. A *monomial walk* on a d -uniform hypergraph is an even sequence of edges $\mathcal{W} := \{e_1, e_2, \dots, e_{2k}\}$, where each edge e_i in the sequence intersects the subsequence $\{e_1, e_2, \dots, e_{i-1}\}$, and each vertex $v \in \bigcup_{i=1}^{2k} e_i$ covered by \mathcal{W} satisfies the following balancing condition: $|\{e_i \in \mathcal{W} : v \in e_i \text{ and } i \text{ is odd}\}| = |\{e_i \in \mathcal{W} : v \in e_i \text{ and } i \text{ is even}\}|$.

Each monomial walk gives rise to a binomial in the ideal I_H . Namely, we say that a binomial $f_{\mathcal{W}}$ arises from \mathcal{W} if

$$f_{\mathcal{W}} := \prod_{i=1}^k t_{e_{2i-1}} - \prod_{i=1}^k t_{e_{2i}}.$$

A monomial walk is said to be *primitive* if there does not exist a proper subwalk, that is, a collection \mathcal{W}' of odd- and even-numbered edges, properly contained in \mathcal{W} , such that \mathcal{W}' is also a monomial walk.

For the remainder of the paper, we occasionally refer to a monomial walk as a walk. For convenience, all hypergraphs in this paper will be uniform. In the non-uniform case, binomials in I_H are not necessarily supported by alternating monomial walks; although one can extend the balancing condition to the non-uniform case. Algebraically, we are simply considering homogeneous toric ideals.

Remark 2.5 (bicoloring the edges in $f_{\mathcal{W}}$). The classical theorems for graphs use odd and even numbered edges to describe $f_{\mathcal{W}}$. For hypergraphs, however, we have found this notation and related pictorial representations extremely tedious. Thus, from now on, we refer to the even- and odd-numbered edges as blue and red edges. This bicoloring is implicit in the graphs case as well.

Example 2.6. Consider a complete 3-uniform hypergraph on 12 vertices. Figure 2 shows three monomial walks on this hypergraph.




FIGURE 2. Examples of Walks

In general, there is no squarefree assumption on the binomial $f_{\mathcal{W}}$. Walking along the same edge multiple times during \mathcal{W} is allowed, so it is possible that $e_i = e_j$ for $i \neq j$. (See Figure 2(C) for an example.) If an edge in a monomial walk appears as both red and blue, then these two copies of the edge can be safely removed. The remaining set of edges constitutes another monomial walk. On the binomial side, this corresponds to a multiple of a smaller binomial.

Definition 2.7. Let \mathcal{E} be a (multiset) collection of an even number of edges in H . We say that \mathcal{E} is *balanced* with respect to a given bicoloring of H if for each vertex v covered by \mathcal{E} , the number of red edges containing v equals the number of blue edges containing v ; in symbols:

$$(\star) \quad \deg_{\text{blue}}(v) = \deg_{\text{red}}(v).$$

The constructions above imply the following.

Theorem 2.8. *Let H be a uniform hypergraph. Any balanced collection of edges $\mathcal{E} \subset E(H)$ constitutes a monomial walk, or a collection of disjoint monomial walks.*

In particular, the ideal I_H is generated by binomials $f_{\mathcal{W}}$ arising from primitive monomial walks \mathcal{W} on H .

Proof. A binomial f is in the ideal I_H if and only if the set of edges \mathcal{E} appearing in the support of the binomial is balanced. Namely, $f \in I_H$ is given by a collection of edges such that

$$f = \prod_{\text{blue edges } e} t_e - \prod_{\text{red edges } e'} t_{e'}$$

where the coloring is induced by the binomial f itself. Being in the kernel of φ_H is equivalent to

$$\varphi_H \left(\prod_{\text{blue edges } e} t_e \right) = \varphi_H \left(\prod_{\text{red edges } e'} t_{e'} \right).$$

This, in turn, is equivalent to Condition (\star) .

We need to show that any balanced collection of edges with respect to some bicoloring can be written as arising from a monomial walk, or a set of disjoint monomial walks. Throughout the proof we will maintain a sequence of visited edges \mathcal{W} from \mathcal{E} , and the set of red and blue degrees of each vertex v covered by \mathcal{E} , which we denote by $\deg_{\text{red}}(v; \mathcal{E})$ and $\deg_{\text{blue}}(v; \mathcal{E})$ respectively.

We begin with $\mathcal{W} = \emptyset$. Pick an arbitrary blue edge $e \in \mathcal{E}$, define $e_1 := e$, and update $\mathcal{W} = \mathcal{W} \cup \{e_1\}$, $\mathcal{E} = \mathcal{E} - \{e_1\}$, noting that $\deg_{\text{blue}}(v; \mathcal{E})$ has decreased by 1, for every $v \in e_1$. Pick an arbitrary $u_1 \in e_1$. By (\star) , there exists a red edge $e' \in \mathcal{E}$ such that $u_1 \in e'$. Let $e_2 := e'$ and update $\mathcal{W} = \mathcal{W} \cup \{e_2\}$, $\mathcal{E} = \mathcal{E} - \{e_2\}$, noting that $\deg_{\text{red}}(v; \mathcal{E})$ has decreased by 1, for every $v \in e_2$. Clearly, u_1 is now balanced in \mathcal{W} .

Repeat recursively as follows. Check whether $\deg_{\text{blue}}(v; \mathcal{E}) = \deg_{\text{red}}(v; \mathcal{E})$ for all v . If not, there must exist an unbalanced vertex covered by \mathcal{W} . For odd (resp. even) i , there exists a blue (resp. red) edge in \mathcal{E} that intersects an even (resp. odd) edge e_j in \mathcal{W} , $j < i$. Call this edge e_i and move it to \mathcal{W} .

On the other hand, if $\deg_{\text{blue}}(v; \mathcal{E}) = \deg_{\text{red}}(v; \mathcal{E})$ holds for every v , we have produced a monomial walk \mathcal{W} . In the case where \mathcal{E} is not empty, the algorithm restarts on the current \mathcal{E} with empty \mathcal{W} to find additional edge-disjoint monomial walks. Finally, the algorithm stops when \mathcal{E} is empty.

What we have shown is that every binomial in the ideal I_H arises from a finite collection of monomial walks. That only primitive binomials are necessary to generate the ideal follows from standard arguments: e.g. see Lemma 4.6 in [Stu96]. \square

Naturally, if a certain edge appeared multiple times in \mathcal{E} , each occurrence of the edge gets a unique label, so that the above argument is valid for multisets of edges.

Recall that the Graver basis is a special generating set of the toric ideal, usually very large and algebraically redundant, but combinatorially nice. What we have shown above is that the Graver basis of I_H consists of primitive monomial walks with support on the edge set of H . In addition, as in the graphs case, since the Graver basis contains all reduced Gröbner bases, we can say the following:

Corollary 2.9. *If f is a polynomial in any reduced Gröbner basis of I_H , then $f = f_{\mathcal{W}}$ for some primitive monomial walk \mathcal{W} .*

The natural problem now is to understand primitive monomial walks on a given hypergraph H . We will concern ourselves with a criterion which will allow us to determine if a given set of edges is a primitive monomial walk. To this end, we make the following definition.

Definition 2.10. A multihypergraph \mathcal{H} is called a *monomial hypergraph* if its edges can be colored so that they are balanced (see Definition 2.7). Such an \mathcal{H} is *primitive* if there exists a balanced bicoloring that makes the corresponding monomial walk primitive.

The purpose of this definition is twofold. First, note that a monomial hypergraph corresponds to a monomial walk. However, a given walk comes equipped with the blue and red bicoloring of the edges (i.e., the binomial is given explicitly); whereas in a monomial hypergraph, the coloring is not fixed (i.e., it represents the support of the binomial). In the graph case, this distinction is less important because the coloring is more obvious. For example, when we say that “even cycles are generators of the toric ideal of a graph”, we do not specify how to color the edges of the cycle; however, upon the inspection of the proof, of course, it is obvious that the edges are colored in alternating colors to give rise to a binomial. In hypergraphs, such a choice of coloring may not always be immediate or unique, so this distinction is important. Second, a monomial hypergraph is not a simple hypergraph: there is no squarefree assumption on the monomial walks, so an edge in a hypergraph may be traced multiple times in a walk over its edge set. This generality is required since we cannot assume that the binomials in I_H are squarefree.

Corollary 2.11. *The set of binomials*

$$\{f_{\mathcal{W}} | \mathcal{W} \text{ is a bicolored primitive monomial hypergraph over the hypergraph } H\}$$

is a universal Gröbner basis for the toric ideal I_H .

In summary: a monomial walk has an alternating coloring on the edges and requires that all the visited vertices are balanced with respect to that coloring; and a monomial hypergraph is a multihypergraph with the supporting edge set being that of H , it corresponds to a monomial walk, but its coloring is not fixed. *The monomial hypergraph corresponding to the monomial walk \mathcal{W} is the combinatorial signature of the binomial $f_{\mathcal{W}} \in I_H$.*

3. MONOMIAL HYPERGRAPHS: FROM LOG-LINEAR MODELS TO DISCREPANCY

Before exploring the structure of monomial hypergraphs in the next section, let us briefly reflect on their relevance. We explain how they arise naturally in algebraic statistics, and proceed to show how they relate to combinatorial discrepancy of hypergraphs.

3.1. Parameter hypergraphs. Monomial maps make a notable appearance in statistics: they encode important families of statistical models for discrete random variables called *log-linear models*. A log-linear model is a family of joint probability distributions that are described by monomials in the model parameters (equivalently, minimal sufficient statistics of the model are in the row span of a linear map, and hence the name; but let us not dwell on this here). To see how monomial maps arise, let us consider the following very simple example. Two random variables Z_1 and Z_2 , taking a and b values, respectively, are said to be independent if the joint probability $P_{ij} := \text{Prob}(Z_1 = i, Z_2 = j)$ factors as follows: $P_{ij} = \text{Prob}(Z_1 = i)\text{Prob}(Z_2 = j)$ for all $i \in [a]$ and $j \in [b]$. Denote the probabilities $\text{Prob}(Z_1 = i)$ and $\text{Prob}(Z_2 = j)$ by x_i and x_{a+j} , respectively. The unknown quantities x_1, \dots, x_{a+b} are called the *model parameters*; they completely determine this family of distributions. We can summarize this simple example by saying that the independence model for Z_1 and Z_2 is specified by the monomial map that sends P_{ij} to $x_i x_{a+j}$. In this sense, the edge subring $k[K_{a,b}]$ of the complete bipartite graph $K_{a,b}$ on the vertex set $\{x_1, \dots, x_a\} \sqcup \{x_{a+1}, \dots, x_{a+b}\}$ describes the structure of the joint probability distributions that belong to the independence model.

This correspondence can be naturally extended to any more complicated model for discrete random variables Z_1, \dots, Z_m , whose joint probabilities are parametrized by squarefree monomials in the model parameters x_1, \dots, x_n . Each such log-linear model \mathcal{M} is encoded by a hypergraph $H_{\mathcal{M}}$ on the vertex set x_1, \dots, x_n , constructed as follows: $\{x_j\}_{j \in J}$ is an edge of $H_{\mathcal{M}}$ if and only if the index set J describes one of the joint probabilities in the model; that is, there exist values i_1, \dots, i_m such that $\text{Prob}(Z_1 = i_1, \dots, Z_m = i_m) = \prod_{j \in J} x_j$.

Henceforth, we will name $H_{\mathcal{M}}$ the *parameter hypergraph of the model \mathcal{M}* . Somewhat surprisingly, toric ideals $I_{H_{\mathcal{M}}}$ are quite relevant for testing goodness of fit of the model \mathcal{M} . Without going into statistical details, a non-asymptotic test for model/data fit is carried out by exploring the fiber of a given point in the image of

the monomial map $\phi_{H_{\mathcal{M}}}$. In terms of the hypergraph $H_{\mathcal{M}}$, an observed data point corresponds to a collection of edges; for convenience, think of them as colored blue. The fiber of that data point then consists of all red edge sets that balance it, and thus have the same image under the map $\phi_{H_{\mathcal{M}}}$. The recipe for exploring a fiber is to move from any point to another point by applying a *Markov move*, which corresponds to a generator of the toric ideal $I_{H_{\mathcal{M}}}$ of the model. Since the binomials in $I_{H_{\mathcal{M}}}$ are supported by balanced edge sets, this move is interpreted as removing the observed (blue) edges and replacing them by the red edges, arriving at another monomial with the same image. The correspondence between Markov moves and toric generators is well-known in algebraic statistics literature, and often referred to as the fundamental theorem of Markov bases. The technical description of how such a random walk on a fiber is actually implemented is beyond the scope of this paper; but the interested reader is referred to [DSS09] and the recent text [AHT12].

In general, finding Markov bases for many relevant statistical models is a wide-open problem; some complexity bounds and a few structural theorems are known. Therefore, our goal is to understand the combinatorial structure of monomial hypergraphs supported on $H_{\mathcal{M}}$. Instead of focusing on one log-linear model, we explore combinatorial signatures of general monomial hypergraphs. We focus on homogeneous ideals, thus the models parametrized by uniform hypergraphs. Finally, as explained in [AHT12], for homogeneous toric models with statistical sampling constraints or with structural zeros in the model, a minimal generating set of the toric ideal may not suffice for connecting the fibers, and instead squarefree part of the Graver basis is used.

As mentioned in the introduction, the edge subring of a *graph* has appeared in several results in algebraic statistics. By defining the parameter hypergraph $H_{\mathcal{M}}$, which the reader should agree is quite a natural generalization, we give an applied motivation to extend the well-known results on toric ideals of graphs to the more general case. By Theorem 2.8, Markov moves for any (homogeneous) log-linear model \mathcal{M} are encoded by primitive monomial hypergraphs supported on a (uniform) $H_{\mathcal{M}}$. One of our main goals is to frame this problem in terms of combinatorics of hypergraphs.

3.2. A dual problem: combinatorial discrepancy. To that end, we begin by noting a relation between monomial hypergraphs and zero-discrepancy hypergraphs. Broadly, the combinatorial discrepancy problem asks for a coloring of the vertices of a hypergraph $H = (V, E)$ such that each hyperedge contains roughly the same number of vertices of each color. (The reader is directed to [Mat10, Cha02] for a detailed introduction to discrepancy theory.) Let us formally define the discrepancy of a hypergraph.

Definition 3.1. The *discrepancy* of a hypergraph $H = (V, E)$, denoted $\text{disc}(H)$, is the quantity $\text{disc}(H) = \min_{\chi: V \rightarrow \pm 1} \max_{e \in E} \sum_{v \in e} \chi(v)$. In words, it is the minimum, over all colorings χ , of the maximum discrepancy of an edge in H .

Given a hypergraph H , the *dual hypergraph* H^* is the hypergraph containing a vertex v^* for each edge $e \in E(H)$ and an edge e^* for each $v \in V(H)$; here $e^* = \{v^* \in V(H^*) : v^* \text{ corresponds to an edge } e \in E(H) \text{ that is incident in } H \text{ to vertex } v\}$. More briefly, the vertex-edge incidence matrix of H^* is the transpose of the vertex-edge incident matrix of H .

It is not hard to see that a zero-discrepancy hypergraph $H^* = (V^*, E^*)$ is dual to a monomial hypergraph $H = (V, E)$: the mapping $\chi : V^* \rightarrow \pm 1$ acts on the vertices of H^* but it acts on the edges of H . Zero discrepancy for H^* implies that each edge of H^* contains the same number of red and blue vertices, and by consequence each vertex of H is incident to the same number of red and blue edges. This establishes the following theorem.

Theorem 3.2. *A hypergraph H is a monomial hypergraph if and only if $\text{disc}(H^*) = 0$. Furthermore, a monomial hypergraph H is primitive if and only if no coloring $\chi : V^* \rightarrow \pm 1$ realizing $\text{disc}(H^*) = 0$, also induces a zero discrepancy coloring on a (non-empty) subhypergraph \tilde{H} of H^* produced by deleting any subset of vertices and shrinking the edges.*

Deleting a vertex from H^* and shrinking its incident edges is equivalent to deleting an edge in the dual hypergraph H . There are some other interesting facts here. While a monomial hypergraph can have multiple edges, this is not true or necessary for the dual zero-discrepancy hypergraph, as multiple edges in \mathcal{H} would correspond to vertices that are incident to the same set of edges in the dual.

An important inapproximability result of [CNN11] on discrepancy implies that it is NP-Hard to decide whether a hypergraph $H = (V, E)$ is the support for a monomial hypergraph, or whether for every bicoloring of $E(H)$, there exists a vertex $v \in V(H)$ such that $|\deg_{\text{red}}(v) - \deg_{\text{blue}}(v)| \geq \Omega(\sqrt{|E(H)|})$.

Given the difficulty of deciding the discrepancy of a hypergraph, we pursue a structural theorem for monomial hypergraphs in the next section.

4. SPARSE BOUQUETS

Theorem 2.8 can be used to give a characterization of (primitive) monomial walks in terms of the existence of an edge partition satisfying the degree condition (\star) . However it gives no other information about the structure of the supporting hypergraphs.

In this section, we use Theorem 2.8 to construct and study supports of monomial hypergraphs that parallel the graphs case. For example, consider case (ii) from Theorem 2.3. Two odd cycles glued at a vertex support a primitive monomial walk on a graph. Increasing the degree of that core vertex could produce monomial walks, but they can never be primitive. In contrast, the core vertices in hypergraph supports of primitive walks can have degree larger than 2; in fact, it can be arbitrarily high!

The first three results in this section directly generalize cases (i) , (ii) and (iii) (for short cycles) from Theorem 2.3. First, a pair of perfect matchings generalizes an even cycle, and its supporting hypergraph is characterized in Proposition 4.2. Second, a monomial sunflower generalizes a bow-tie (two 3-cycles sharing a vertex), and its supporting hypergraphs are characterized in Proposition 4.5. Finally, a partitioned core sunflower generalizes two 3-cycles connected by two paths, and the supporting hypergraphs are described in Proposition 4.9.

A further generalization summarizing the above propositions is given in Theorem 4.12. So, unlike in the case of graphs, these three results do not exhaustively characterize monomial hypergraphs. Nevertheless, starting from sunflowers and matchings, natural building blocks of monomial hypergraphs, we can obtain much more complex structures. We delve more into the open problem of structurally characterizing primitive monomial hypergraphs in Section 6.

4.1. Matchings.

Definition 4.1. A *matching* on a hypergraph $H = (V, E)$ is a subset $M \subset E$ of independent edges, that is, no two edges intersect. A matching is called *perfect* if it covers all the vertices of the hypergraph, i.e., $V(M) = V$.

The simplest monomial hypergraph can be formed using a pair of perfect matchings on the same set of variables. Next we give a sufficient and necessary condition for the primitivity of such a monomial hypergraph. A hypergraph $H = (V, E)$ is said to be *connected* if its primal graph is connected, where the primal graph has the same vertex set and an edge between any two vertices contained in the same hyperedge.

Proposition 4.2 (Pair of matchings). *Let $H = (V, E)$ be a d -uniform hypergraph such that $E = M_r \sqcup M_b$, where M_r, M_b form two edge-disjoint perfect matchings on the vertex set V . A monomial hypergraph \mathcal{H} with support H is primitive if and only if H is connected and \mathcal{H} contains no multiple edges.*

Proof. Let \mathcal{H} be a primitive monomial hypergraph and consider a primitive monomial walk \mathcal{W} on the edges of \mathcal{H} with an appropriate associated bicoloring. Any walk is connected, so suppose for contradiction that edge $e \in E$ has multiplicity k in \mathcal{W} . By construction of the supporting hypergraph H , there exist exactly two distinct edges in H , say e and f , containing any vertex v . We may assume that all copies of e in \mathcal{W} are red, as two copies of e in \mathcal{W} that belong to a different color partition would create a trivial subwalk. By the degree condition (\star) there must be k copies of f in \mathcal{W} , all colored blue. Edges e and f are distinct so there exists a vertex u in f and a unique edge g in H such that g contains u and has multiplicity k in \mathcal{W} . But H is connected, so we may repeat this process for each vertex in \mathcal{W} to conclude that all edges in E appear in \mathcal{W} in k copies of the same color, contradicting the primitivity of \mathcal{W} .

For sufficiency, we need only produce a bicoloring for the edges of H , taken with no multiplicity, such that the associated walk is primitive. We construct the required bicoloring by coloring edges in M_r red and edges in M_b blue. Assume by way of contradiction that there exists a smaller walk \mathcal{W} in \mathcal{H} that is also a monomial walk. Let $V_{\mathcal{W}}$ and $E_{\mathcal{W}}$ be the vertex and edge sets induced by the walk \mathcal{W} , and let $\mathcal{H}_{\mathcal{W}} = (V_{\mathcal{W}}, E_{\mathcal{W}})$. It is clear that every vertex in any monomial walk, and so specifically in $\mathcal{H}_{\mathcal{W}}$ must have degree at least 2 (the sum of the blue degree and the red degree). Since \mathcal{W} is a subwalk, there must be some edge or vertex of \mathcal{H}

that \mathcal{W} does not contain. But every vertex in $H_{\mathcal{W}}$ has degree 2 so there can be no more edges in $\mathcal{H} \setminus \mathcal{H}_{\mathcal{W}}$ containing any vertex from $\mathcal{H}_{\mathcal{W}}$. Then H is disconnected. \square

Figures 2(A) and 2(B) are primitive matchings. The former is also commonly called a *tight cycle*.

4.2. Monomial Sunflower. As pairs of matchings have the property that each vertex has degree 2, they generalize even cycles on graphs. In contrast, the next case in Theorem 2.3 has a distinguished vertex of higher degree. Its natural generalization in hypergraphs, which we will call a monomial sunflower, is the topic of this subsection, and is based on a very useful, highly structured hypergraph, the *sunflower*, which, incidentally, is guaranteed to occur in hypergraphs with large enough edge sets, independently of the size of the vertex set. (See e.g. [Juk01].)

Definition 4.3. A d -uniform hypergraph $H = (V, E)$ is a *sunflower* if $e_i \cap e_j = C$ for all edges $e_i \neq e_j \in E$, and $C \subset V$. The set of vertices C is called the *core* of the sunflower, and each e_i is called a *petal*.

In general, one may allow the core to be empty, in which case you simply get a matching. In this paper, we will explicitly use the name matching instead of allowing for empty cores.

For the remainder of the paper, fix the following notation.

Notation 4.4. For a subset $C \subset V(H)$, define $H - C$ to be the hypergraph with vertex set $(V \setminus C)$ and edges $\{e_i \setminus C : e_i \in E(H)\}$. Similarly, for a subset $C \subset V(\mathcal{H})$, define $\mathcal{H} - C$ to be the multihypergraph with vertex set $(V \setminus C)$ and edges $\{e_i \setminus C : e_i \in E(\mathcal{H})\}$. The hypergraphs $H - C$ and $\mathcal{H} - C$ are not necessarily uniform and may include empty edges, or multiple copies of an edge. See Figures 3, 4 and 9 for examples.

Clearly, there is no nontrivial monomial walk on a sunflower: since the degree of any non-core vertex is 1, every petal needs to be both blue and red, resulting in the zero binomial. Thus, a sunflower alone cannot be a support of a monomial hypergraph. The non-core petal vertices could, however, be balanced by a perfect matching. To that end, consider H to be a sunflower with a perfect matching on its non-core vertices. In Proposition 4.5, we show that a primitive monomial hypergraph with support H can be characterized by counting petals.

Consider the set of connected components

$$H - C = \bigcup_{j \in \mathcal{I}} G_j.$$

For convenience, define $G_{\mathcal{J}} := \bigcup_{j \in \mathcal{J}} G_j$, for any index set $\mathcal{J} \subseteq \mathcal{I}$. By abuse of notation we will write $|G_j|$ for the number of edges in the multiset $\{e \in H : e \cap C \neq \emptyset \text{ and } e \setminus C \in G_j\}$.

The next proposition is based on the idea that counting petals detects balanced walks on H .

Proposition 4.5 (Sunflower). *Let H be a hypergraph consisting of a sunflower and a perfect matching on the non-core vertices of the sunflower. The hypergraph H is the support of a monomial hypergraph \mathcal{H} if and only if there exists a partition $(G_{\mathcal{J}}, G_{\mathcal{K}})$ of the connected components of $H - C$ such that*

$$\sum_{j \in \mathcal{J}} m_j |G_j| = \sum_{k \in \mathcal{K}} m_k |G_k|, \quad \text{where } m_j, m_k \text{ are integers.}$$

In addition, \mathcal{H} is primitive if and only if the partition $(G_{\mathcal{J}}, G_{\mathcal{K}})$ also satisfies the following two conditions:

- i. *no two subsets $\mathcal{J}' \subsetneq \mathcal{J}$ and $\mathcal{K}' \subsetneq \mathcal{K}$ satisfy $\sum_{j \in \mathcal{J}'} m'_j |G_j| = \sum_{k \in \mathcal{K}'} m'_k |G_k|$ for some integers $m'_j \leq m_j$, $m'_k \leq m_k$ and*
- ii. *there are no integers $m'_j, m'_k \in \mathbb{Z}$ such that the partition $(G_{\mathcal{J}}, G_{\mathcal{K}})$ of the connected components of $H - C$ satisfies $\sum_{j \in \mathcal{J}} m'_j |G_j| = \sum_{k \in \mathcal{K}} m'_k |G_k|$, where $m'_j \leq m_j$, $m'_k \leq m_k$ and for at least one integer the inequality is strict.*

Proof. For the easy direction of the first statement, suppose $(G_{\mathcal{J}}, G_{\mathcal{K}})$ is a partition of the connected components of $H - C$ such that $\sum_{j \in \mathcal{J}} m_j |G_j| = \sum_{k \in \mathcal{K}} m_k |G_k|$ for some $m_j, m_k \in \mathbb{Z}$. Construct the monomial walk \mathcal{W} as follows: for every edge e having a nonempty intersection with $G_j, j \in \mathcal{J}$ (respectively $G_k, k \in \mathcal{K}$), include in \mathcal{W} , m_j (respectively m_k) copies of edge e . Color the matching edges red for \mathcal{J} and blue for \mathcal{K} , and

use the opposite colors for the petal edges. This immediately ensures that the non-core vertices satisfy the degree condition (\star) . The degrees of a core vertex v are $\deg_{\text{blue}}(v) = \sum_{j \in \mathcal{J}} m_j |G_j|$ and $\deg_{\text{red}} = \sum_{k \in \mathcal{K}} m_k |G_k|$ respectively.

For the other direction of the first statement, suppose \mathcal{H} is a monomial hypergraph over H and consider the monomial walk \mathcal{W} consisting of the edges of \mathcal{H} . For every vertex v in a component G_i of H , there are exactly two distinct edges containing v : a petal edge e and a matching edge f . Suppose e appears with multiplicity m_i in \mathcal{W} . Without loss of generality, we may assume that all m_i copies of e must belong to the same color partition, say red. By the degree condition (\star) , edge f must also appear in \mathcal{W} with multiplicity m_i , and all m_i copies must belong to the blue color partition. Now there exists a vertex $u \in f \cap e'$ ($u \neq v$) where e' ($e' \neq e$) is the unique petal edge containing u . To ensure u is balanced, by the degree condition, e' must also appear in \mathcal{W} with multiplicity m_i . Repeating this argument over all vertices in the connected component G_i we conclude that each matching edge in G_i and each petal edge whose non-core vertices fall in G_i , must appear in \mathcal{W} with multiplicity m_i . Then we may construct the partition $(G_{\mathcal{J}}, G_{\mathcal{K}})$ so that $G_i \in G_{\mathcal{J}}$ if the matching edges it contains are colored red, and $G_i \in G_{\mathcal{K}}$ otherwise. Since the core vertices are also balanced, the equation $\sum_{j \in \mathcal{J}} m_j |G_j| = \sum_{k \in \mathcal{K}} m_k |G_k|$ holds, where the sum on the left hand side counts the blue degree of a core vertex and the right hand side its red degree.

The second statement in the proposition (primitivity) follows from the first statement and the definition of primitivity. Condition *i.* describes all subwalks not using all edges of the supporting hypergraph H and condition *ii.* describes all subwalks using all edges of H with smaller multiplicity than in \mathcal{H} . \square

We may now make the following definition.

Definition 4.6. A *monomial sunflower* is a monomial hypergraph with support set consisting of a sunflower and a perfect matching on the non-core vertices.

The partition $(G_{\mathcal{J}}, G_{\mathcal{K}})$ corresponds to a bicoloring of the petals of the sunflower, as illustrated in the following example.

Example 4.7. Figure 3 demonstrates that H_1 is the support of a simple non-primitive monomial sunflower and its components. Each edge appears exactly once in H_1 , and it is easy to see that for any two components in $H - C$, there is a bicoloring of \mathcal{H}_1 such that the two components form a subwalk.




FIGURE 3. A walk on the non-primitive monomial sunflower \mathcal{H}_1 (left) and the components of $H_1 - C$ with appropriate coloring and multiplicity (right).

Figure 4 demonstrates a bicoloring of the monomial sunflower \mathcal{H}_2 that makes it primitive. The appropriate multiplicities for each component are $m_1 = 1$ and $m_2 = 2$.

4.3. Independent Petal Bouquets. To produce more general monomial hypergraphs and characterize primitivity on them, we may generalize the monomial sunflower by considering a hypergraph H formed by a collection of sunflowers and a perfect matching on the subset of vertices that are not in the core of any sunflower. Let us, first, assume that the sunflowers are entirely disjoint.




FIGURE 4. A walk on the primitive monomial sunflower \mathcal{H}_2 (left) and the components of $H_1 - C$ with appropriate bicoloring and multiplicity (right).

Let H be a hypergraph consisting of a collection of vertex-disjoint sunflowers S_1, S_2, \dots, S_ℓ and a perfect matching on the non-core vertices of $\bigcup_i S_i$. We will call such an H a *matched-petal partitioned-core sunflower*.

As we saw in Proposition 4.5, when H contains a single sunflower, the components of $H - C$ and their multiplicities determine the degrees of the vertices of H and ensure that the core is balanced. In a matched-petal partitioned-core sunflower, not all the components contribute to every sunflower S_i . So, one needs to keep track of the components of $H - C$ that contain vertices of S_i , for *each* S_i . We employ the following notation.

Notation 4.8. Let H be a matched-petal partitioned-core sunflower. Given a fixed partition $(G_{\mathcal{J}}, G_{\mathcal{K}})$ of the connected components of $H - C$, and a fixed sunflower S_i of H , define $\mathcal{J}(S_i) \subset \mathcal{J}$ (respectively $\mathcal{K}(S_i) \subset \mathcal{K}$) to include every $j \in \mathcal{J}$ (respectively $k \in \mathcal{K}$), such that G_j contains a vertex of S_i . In symbols, $\mathcal{J}(S_i)$ is the set of indices defined as $\mathcal{J}(S_i) := \{j \in \mathcal{J}(S_i) : j \in \mathcal{J} \text{ and } e \cap G_j \neq \emptyset \text{ for some edge } e \in S_i\}$.

Proposition 4.9 (Partitioned Core Sunflower). *Let H be a matched-petal partitioned-core sunflower, and adopt notation 4.8.*

The hypergraph H is the support of a monomial hypergraph \mathcal{H} if and only if there exists a partition $(G_{\mathcal{J}}, G_{\mathcal{K}})$ of the connected components of $H - C$ such that for each $1 \leq i \leq \ell$:

$$\sum_{j \in \mathcal{J}(S_i)} m_j |G_j| = \sum_{k \in \mathcal{K}(S_i)} m_k |G_k|, \quad \text{where } m_j, m_k \text{ are integers.}$$

In addition, \mathcal{H} is primitive if and only if the partition $(G_{\mathcal{J}}, G_{\mathcal{K}})$ also satisfies the following two conditions:

- i. *no two subsets $\mathcal{J}' \subsetneq \mathcal{J}$ and $\mathcal{K}' \subsetneq \mathcal{K}$ satisfy for each $1 \leq i \leq \ell$,*

$$\sum_{j \in \mathcal{J}(S_i) \cap \mathcal{J}'} m'_j |G_j| = \sum_{k \in \mathcal{K}(S_i) \cap \mathcal{K}'} m'_k |G_k|,$$

for some integers $m'_j \leq m_j, m'_k \leq m_k$, and

- ii. *there are no integers $m'_j, m'_k \in \mathbb{Z}$ such that the partition $(G_{\mathcal{J}}, G_{\mathcal{K}})$ of the connected components of $H - C$ satisfies for each $1 \leq i \leq \ell$,*

$$\sum_{j \in \mathcal{J}(S_i)} m'_j |G_j| = \sum_{k \in \mathcal{K}(S_i)} m'_k |G_k|,$$

where $m'_j \leq m_j, m'_k \leq m_k$ and for at least one integer the inequality is strict.

Proof. The first statement follows with the same argument as Proposition 4.5 for the non-core vertices. For a core vertex, given a partition $(G_{\mathcal{J}}, G_{\mathcal{K}})$ of the connected components of $H - C$, the sum $\sum_{j \in \mathcal{J}(S_i)} m_j |G_j|$ (respectively $\sum_{k \in \mathcal{K}(S_i)} m_k |G_k|$) counts the number of petals of sunflower S_i whose remnants fall in components $G_{\mathcal{J}(S_i)}$ (respectively $G_{\mathcal{K}(S_i)}$), and as such completely determines the quantity that each part of the

bipartition contributes to the degree of a core vertex $v \in S_i$. As before, the second statement describes two types of subwalks. \square

Example 4.10. Figure 5 demonstrates a primitive partitioned core monomial sunflower \mathcal{H} . Note that all of the multiplicities are equal to one.




FIGURE 5. Primitive partitioned-core monomial sunflower \mathcal{H}_3 , with an appropriate bicoloring.

We relax the disjointness property of the sunflowers above, and allow the sunflowers to intersect on core vertices. Let H be a hypergraph consisting of a collection of sunflowers S_1, S_2, \dots, S_ℓ , which may only intersect at their cores (known as a relaxed core sunflower), and a perfect matching on the non-core vertices of the sunflowers. We will call such an H a *matched-petal relaxed-core sunflower*.

In the partitioned core case, the degree of each core vertex depends on exactly one sunflower. In contrast, allowing intersections in the cores causes the degree of a core vertex to depend on several sunflowers. This prompts the following notation.

Notation 4.11. For a matched-petal relaxed-core sunflower, define $\mathcal{I}(v)$ to consist of the set of indices i such that vertex v is in the core of sunflower S_i . In symbols $\mathcal{I}(v) := \{i \in \mathcal{I}(v) : v \in S_i\}$.

In addition, notation 4.8 extends to H .

Theorem 4.12 (Relaxed Core Sunflower). *Let H be a matched-petal relaxed-core sunflower, and adopt notation 4.8 and 4.11.*

The hypergraph H is the support of a monomial hypergraph \mathcal{H} if and only if there exists a partition $(G_{\mathcal{J}}, G_{\mathcal{K}})$ of the connected components of $H - C$ such that for each vertex v in the core of a sunflower:

$$\sum_{i \in \mathcal{I}(v)} \sum_{j \in \mathcal{J}(S_i)} m_j |G_j| = \sum_{i \in \mathcal{I}(v)} \sum_{k \in \mathcal{K}(S_i)} m_k |G_k|, \quad \text{where } m_j, m_k \text{ are integers.}$$

In addition, \mathcal{H} is primitive if and only if the partition $(G_{\mathcal{J}}, G_{\mathcal{K}})$ also satisfies the following two conditions:

- i. no two subsets $\mathcal{J}' \subsetneq \mathcal{J}$ and $\mathcal{K}' \subsetneq \mathcal{K}$ satisfy for each vertex v in the core of the sunflower,*

$$\sum_{i \in \mathcal{I}(v)} \sum_{j \in \mathcal{J}(S_i) \cap \mathcal{J}'} m'_j |G_j| = \sum_{s \in \mathcal{I}(v)} \sum_{k \in \mathcal{K}(S_i) \cap \mathcal{K}'} m'_k |G_k|,$$

for some integers $m'_j \leq m_j$, $m'_k \leq m_k$, and

- ii. there are no integers $m'_j \leq m_j$, $m'_k \leq m_k \in \mathbb{Z}$, where the inequality is strict for at least one integer, such that the partition $(G_{\mathcal{J}}, G_{\mathcal{K}})$ of the connected components of $H - C$ satisfies for each vertex v in the core of a sunflower,

$$\sum_{i \in \mathcal{I}(v)} \sum_{j \in \mathcal{J}(S_i)} m'_j |G_j| = \sum_{i \in \mathcal{I}(v)} \sum_{k \in \mathcal{K}(S_i)} m'_k |G_k|.$$

Proof. Follow the proof of Proposition 4.9 with one important exception. The degree of a core vertex may now depend on several sunflowers, and so the quantity necessary to calculate the degree of a core vertex v in part \mathcal{J} , for example, is $\sum_{i \in \mathcal{I}(v)} \sum_{j \in \mathcal{J}(S_i)} m'_j |G_j|$. \square

Proposition 4.9 can be obtained as a corollary of Theorem 4.12.

Example 4.13. Figure 6 demonstrates a simple primitive relaxed core monomial sunflower \mathcal{H}_4 with four sunflowers. Each multiplicity equals one.




FIGURE 6. Primitive relaxed core monomial sunflower \mathcal{H}_4 , with an appropriate bicoloring.

5. BOUQUET COMPLEXITY

The types of Graver elements studied in Propositions 4.2, 4.5 and 4.9 are, in fact, the *simplest* ones that may occur beyond the graphs case. They are the natural direct generalization of the three types of monomial walks on graphs from Theorem 2.3 (see also Figure 1), each of which is an example of a sparse bouquet. By allowing the edges to contain more than two vertices, we inadvertently and significantly increase the complexity of the primitive walks. This already occurs on 3-uniform hypergraphs and on sparse bouquets. We summarize some of the different ways in which the complexity has increased through the following examples.

In the case of graphs, it is known (see [Vil00, Proposition 8.1.6]) that the largest degree of a vertex in the support of a primitive walk is 4, and that the edges in the support of the walk can be traversed at most twice. None of these restrictions extend to d -uniform hypergraphs ($d > 2$), where a primitive walk can contain vertices with arbitrarily large degree; can contain arbitrarily many vertices of large degree; and a supporting edge can be used arbitrarily many times in a primitive walk.

It is not hard to see that there exists exactly one primitive d -uniform monomial sunflower with core of size $d - 1$: the one containing exactly $2d + 2$ hyperedges. On the other hand, figure 7 demonstrates a family of primitive 3-uniform bouquets on n vertices with number of edges of the order of n and containing a vertex with degree of the order of n . Figure 8 demonstrates a family of 3-uniform bouquets in which the number of cores, vertex degrees, and edge multiplicities can be arbitrarily high.




FIGURE 7. A primitive monomial hypergraph containing a vertex with arbitrarily large degree.




FIGURE 8. A primitive monomial hypergraph with arbitrarily many cores.

6. CONCLUSION AND OPEN PROBLEMS

Using the sparse hypergraphs as a starting point, one should consider the problem of classifying supports of more general monomial hypergraphs. However, this problem is highly nontrivial. For example, Figure 9 represents a simplest example of a non primitive monomial walk that is not of the form covered in Section 4. The hypergraph consists of two overlapped walks. It is tempting to mimic the constructions for sparse bouquets and express this hypergraph as a union of, say, four sunflowers S_1, \dots, S_4 with two petals each: $S_1 = \{e_1, e_3\}$, $S_2 = \{e_2, e_4\}$, $S_3 = \{e_5, e_7\}$, $S_4 = \{e_6, e_8\}$. Then, removing the cores from these sunflowers and counting the petal remnants in the connected components, with multiplicity 1, as suggested by a direct generalization of Theorem 4.12, would reveal that $S_3 \cup S_4$ is a balanced set of edges. In this way, one could detect that the monomial hypergraph $S_1 \cup \dots \cup S_4$ is not primitive.

In general, however, an arbitrary monomial hypergraph consists of sunflowers that overlap in unrestricted ways, and thus it is not guaranteed that removing the cores will disconnect and detect a subwalk. In particular, a direct generalization of the results from Section 4 does not hold without further assumptions.

Problem 6.1. Find a criterion to detect primitivity of a given monomial hypergraph H based on a subset of the edges in H .




FIGURE 9. A non-primitive monomial bouquet

At the moment, the best result in this direction is Theorem 2.8. For “sparse” collections of sunflowers and matchings, Section 4 offers more efficient criteria. But if the hypergraph is more dense, an arbitrary sunflower decomposition alone does not seem to suffice.

More generally, it would be interesting, although difficult, to solve the following specific instance of the general problem of characterizing primitive walks.

Problem 6.2. Characterize (combinatorially) the supports of primitive monomial hypergraphs over an arbitrary 3-uniform hypergraph.

We remind the reader that even for 3-uniform hypergraphs, any of the types of bouquets that we describe can appear in the Graver basis, as demonstrated throughout the paper. Thus Problem 6.2 is highly non-trivial; however, we expect it to have a nice answer for a subclass of 3-uniform hypergraphs, for example, where it is possible to partition a vertex set in a way that restricts the types of bouquets that may appear. On the other hand, recall the ultimate goal of interest to statistical applications: we are looking for characterizations of the structure of monomial walks for a family of hypergraphs specified by conditions that are significantly more restrictive than uniformity. The family will come from a statistical model with a prescribed edge size, or edge types. Such conditions will necessarily restrict the ways in which edges of H can intersect, and thus may imply the sparsity of walks supported on H .

Furthermore, one can quickly derive the entire Graver basis for the following two classes of hypergraphs. For *any* 2-regular hypergraph, it consists only of pairs of matchings as in Proposition 4.2. In particular, the Graver basis of a connected 2-regular hypergraph contains exactly one element. Secondly, consider the d -uniform hypergraph $C_r^d = (V, E)$, produced by letting $V = \bigsqcup_{1 \leq i \leq r} V_i$, with each V_i having size $d/2$ (d even), and edge set $E = \{V_i \cup V_j : \text{for } 1 \leq i < j \leq r\}$. A (primitive) monomial walk on C_r^d corresponds to a (primitive) even closed walk in the complete graph K_r , as each hyperedge in C_r^d is in a one-to-one correspondence with an edge in K_r .

Our next problem concerns generating primitive monomial hypergraphs. It is possible to generate such hypergraphs by contraction or “gluing” operations using known monomial hypergraphs (and sparse bouquets in particular) as building blocks. For example, one may produce d -uniform monomial sunflowers from a d -uniform pair of perfect matchings by identifying selected sets of independent vertices. This operation preserves primitivity. The converse is only true if vertex-identification takes place within connected components. In general, one may also glue two hypergraphs H_1, H_2 along the vertices of an edge to produce a new hypergraph H , introducing new primitive elements in I_H constructed from binomials from I_{H_1} and I_{H_2} . The support of the monomial hypergraph in Figure 10 was partially constructed from the support of four monomial sunflowers with this gluing operation. It is not hard to see that \mathcal{B} is a monomial hypergraph.

A general gluing operation for graphs, based on a toric fiber product, is described in [EKS]. It preserves many properties of toric ideals.

Problem 6.3. Find a gluing operation on hypergraphs that preserves primitivity of monomial walks. Is it possible to generate all primitive monomial hypergraphs starting from sparse bouquets?




FIGURE 10. A primitive bouquet \mathcal{B} partially produced by gluing four sparse bouquets.

The primitive walks in hypergraphs are clearly much more general than in the case of graphs, as Section 5 shows. Nevertheless, we expect many of the other properties of the coordinate ring of I_H to have combinatorial interpretations. For example, the first author studies the degree bounds on the minimal generators of I_H in [GP], generalizing the theorems known for graphs.

Finally, the problems of finding nice term orders and reduced Gröbner bases of I_H , characterizing Cohen-Macaulayness and normality of the coordinate ring and the corresponding polytope, and relating known coloring-inspired properties of hypergraphs to various invariants of the toric ideal I_H are wide open.

ACKNOWLEDGEMENTS

The authors thank György Turán for inspiring discussions, at the start of this project, about the fundamental problem of walks in hypergraphs; and also Elizabeth Gross for suggesting the name “bouquet”, and proofreading an early version of the manuscript. The second author is grateful to Amitava Bhattacharya for introducing her to combinatorial discrepancy. Finally, we are grateful to the anonymous referee for thoughtful and thorough comments that helped improve this work.

The authors acknowledge support by grant FA9550-12-1-0392 from the U.S. Air Force Office of Scientific Research (AFOSR) and the Defense Advanced Research Projects Agency (DARPA). The AMS Simons Travel Grant has helped the first author tremendously during the first year of the project.

REFERENCES

- [AHT12] Satoshi Aoki, Hisayuki Hara, and Akimichi Takemura, *Markov bases in algebraic statistics*, Springer Series in Statistics, Springer New York, 2012.
- [Cha02] Bernard Chazelle, *The discrepancy method: Randomness and complexity*, Cambridge University Press, 2002.
- [CN09] Alberto Corso and Uwe Nagel, *Monomial and toric ideals associated to Ferrers graphs*, Trans. Amer. Math. Soc. **361** (2009), no. 1371–1395.
- [CNN11] Moses Charikar, Alantha Newman, and Aleksandar Nikolov, *Tight hardness results for minimizing discrepancy*, Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '11, 2011, pp. 1607–1614.
- [DSS09] Mathias Drton, Bernd Sturmfels, and Seth Sullivant, *Lectures on algebraic statistics*, Oberwolfach Seminars, vol. 39, Birkhäuser, 2009.
- [DV11] Luis A. Dupont and Rafael H. Villarreal, *Algebraic and combinatorial properties of ideals and algebras of uniform clutters of TDI systems*, Journal of Combinatorial Optimization **21** (2011), no. 3, 269–292.
- [EKS] Alexander Engstrom, Thomas Kahle, and Seth Sullivant, *Multigraded commutative algebra of graph decompositions*, arXiv:1102.2601.
- [GP] Elizabeth Gross and Sonja Petrović, *Combinatorial degree bound for toric ideals of hypergraphs*, arXiv:1206.2512.
- [HT08] Huy Tài Hà and Adam Tuyl, *Monomial ideals, edge ideals of hypergraphs, and their graded Betti numbers*, Journal of Algebraic Combinatorics: An International Journal **27** (2008), no. 2, 215–245.
- [Juk01] Stasys Jukna, *Extremal combinatorics - with applications in computer science*, Texts in theoretical computer science, Springer, 2001.
- [Mat10] Jiri Matoušek, *Geometric discrepancy: An illustrated guide*, Springer Verlag, 2010.

- [Mor13] Jason Morton, *Relations among conditional probabilities*, Journal of Symbolic Computation **50** (2013), 478–492.
- [OH99] Hidefumi Ohsugi and Takayuki Hibi, *Toric ideals generated by quadratic binomials*, Journal of Algebra **218** (1999), no. 2, 509–527.
- [OH00] ———, *Compressed polytopes, initial ideals and complete multipartite graphs*, Illinois J of Mathematics **44** (2000), no. 2, 391–406.
- [OHT13] Mitsunori Ogawa, Hisayuki Hara, and Akimichi Takemura, *Graver basis for an undirected graph and its application to testing the beta model of random graphs*, Annals of Institute of Statistical Mathematics. **65** (2013), 191–212.
- [PRF10] Sonja Petrović, Alessandro Rinaldo, and Stephen E. Fienberg, *Algebraic statistics for a directed random graph model with reciprocation*, Algebraic Methods in Statistics and Probability II (M. Viana and H. Wynn, eds.), Contemporary Mathematics, vol. 516, American Mathematical Society, Providence RI, 2010, pp. 261–283.
- [RPF] Alessandro Rinaldo, Sonja Petrović, and Stephen E. Fienberg, *Maximum likelihood estimation in the β model*, Annals of Statistics. To appear. Preprint available at arXiv:1105.6145.
- [RTT12] Enrique Reyes, Christos Tatakis, and Apostolos Thoma, *Minimal generators of toric ideals of graphs*, Advances in Applied Mathematics **48** (2012), no. 1, 64–78.
- [Stu96] Bernd Sturmfels, *Gröbner bases and convex polytopes*, University Lecture Series, no. 8, American Mathematical Society, 1996.
- [SVV94] Aron Simis, Wolmer V. Vasconcelos, and Rafael H. Villarreal, *On the ideal theory of graphs*, Journal of Algebra **167** (1994), no. 2, 389–416.
- [SW12] Bernd Sturmfels and Volkmar Welker, *Commutative algebra of statistical ranking*, Journal of Algebra **361** (2012), 264–286.
- [Vil95] Rafael H. Villarreal, *Rees algebras of edge ideals*, Communications in Algebra **23** (1995), no. 9, 3513–3524.
- [Vil00] Rafael H. Villarreal, *Monomial algebras*, CRC Press, 2000.

SONJA PETROVIĆ IS WITH THE DEPARTMENT OF STATISTICS, THE PENNSYLVANIA STATE UNIVERSITY, UNIVERSITY PARK PA 16802
E-mail address: petrovic@psu.edu

DESPINA STASI WAS WITH THE DEPARTMENT OF MATHEMATICS, STATISTICS AND COMPUTER SCIENCE, UNIVERSITY OF ILLINOIS AT CHICAGO, CHICAGO IL 60607; AND IS CURRENTLY AT THE DEPARTMENT OF STATISTICS, THE PENNSYLVANIA STATE UNIVERSITY, UNIVERSITY PARK PA 16802
E-mail address: despina@psu.edu